E-ISSN 2586-9027
Homepage : *https://tci-thaijo.org/index.php/SciTechAsia* Science & Technology Asia
Vol.30 No.3 July - September 2025 Page: [143-153]

Original research article

A Preliminary Exploration of Energy Consumption in Blockchain Applications for Thai Healthcare Networks

Supakit Prueksaaroon*, Supakrit Nithikethkul

Department of Electrical and Computer Engineering, Faculty of Engineering, Thammasat School of Engineering, Thammasat University, Pathum Thani 12120, Thailand

Received 12 June 2025; Received in revised form 24 August 2025
Accepted 27 August 2025; Available online 30 September 2025

ABSTRACT

Thailand's fragmented healthcare information systems limit interoperability, seamless data exchange, and nationwide analytics. The Thai Health Data Center uses the "43-files standard" to aggregate hospital data through 13 regional centers before reaching the central database. However, this centralized model struggles with availability and resilience—critical for sensitive, time-critical healthcare data tied to legal and insurance processes. Ensuring data security, accessibility, and protection against loss is vital, especially as expectations for disaster recovery grow amid risks like outages and natural disasters. Blockchain offers a secure and scalable alternative for managing healthcare data. In national systems where trust exists among participants, permissioned blockchain is suitable, providing controlled access, stronger security, and greater availability. To limit energy use, the proposed design employs Proof of Authority (PoA), a consensus mechanism optimized for trusted environments. Using CloudSim Plus, the study simulates and compares energy consumption between classical and blockchain-based systems. The architecture comprises 23 authorized nodes—including regional health offices, zoning divisions, insurance agencies, and the Ministry of Public Health. Transactions are modeled as cloudlets executed on Dell PowerEdge XR11 servers, with volume estimated via regression from 2018–2023 data at about 75 million monthly transactions. Results show annual energy use of 15,827.62 kWh for classical systems and 29,410.02 kWh for blockchain, reflecting a 1.86-fold increase due to added resilience and continuity.

Keywords: Blockchain; CloudSim plus; Energy consumption; Proof of authority; Thai healthcare network

***Corresponding author:** psupakit@engr.tu.ac.th **doi:** 10.14456/scitechasia.2025.57

1. Introduction

In Thailand, hospitals operate under different administrative authorities and often use incompatible health information systems, which hinders the seamless exchange of medical records. Separate ministries govern public hospitals, while private hospitals function independently, resulting in a fragmented healthcare network. Due to system incompatibility and privacy regulations, patient data cannot be directly shared between institutions. Consequently, patients often must retrieve and deliver their own records when moving between facilities, which can cause delays in urgent care. Incomplete or missing data may lead to misdiagnosis, redundant procedures, or inappropriate treatment, compromising patient safety. A similar issue exists with insurance claims, where the lack of connectivity between hospitals and insurers results in manual, delayed reimbursements.

To support health data collection and reporting, the Thai healthcare system is regionally governed through the National Health Security Office (NHSO), which administers healthcare services across 13 health districts [1]. In collaboration with the Thai Health Data Center, the NHSO introduced the "43-files standard" [2], a structured framework for transmitting data from hospitals and clinics to regional and national data centers. This facilitates standardized data exchange, enabling nationwide analytics, epidemic modeling, and health policy development. However, the 43-files system remains centralized and relies on manual or semi-automated submissions. Real-time interoperability between disparate hospital systems is still lacking, underscoring the need for a more integrated and secure solution.

To address the limitations of Thailand's current centralized healthcare data infrastructure, this study explores the adoption of blockchain technology for managing health information. Centralized systems face several critical challenges:

- Scalability constraints: All operations must pass through a single point of processing.
- Single points of failure: Disruption at any central server can halt data access.
- Risk of data tampering: Data can be overwritten or altered, compromising the integrity of medical records.

The Blockchain Trilemma is a conceptual limitation in blockchain technology, where only two of the three key properties—scalability, security, and decentralization—can be fully achieved at a time [3]. Removing trust between organizations is not feasible due to existing policies and would require significant changes. Therefore, the blockchain system can be designed to prioritize scalability and security, while leaving decentralization aside, without challenging the Blockchain Trilemma. This approach offers several benefits, including efficient data exchange across a large-scale network while ensuring robust security, making it well-suited for the complex and highly regulated healthcare environment. Key advantages include [4]:

1. Redundancy and fault tolerance: Ensures data remains accessible even if some nodes fail.
2. Peer-to-peer communication – Reduces reliance on central infrastructure and minimizes bottlenecks.

3. Immutability – Guarantees data integrity by making it append-only, ensuring a secure and auditable history.

Moreover, the blockchain type will be permissioned, ensuring that only predefined validators are allowed to participate in the consensus process [5]. However, blockchain's energy consumption is a well-known concern [6-9]. This study focuses on evaluating the trade-offs by estimating the energy usage of a blockchain-based architecture in comparison to a centralized system, using CloudSim Plus as the simulation framework [10].

The objectives of this study are to review Thailand's existing healthcare data governance framework in order to inform the design of a system that aligns with current practices and legacy infrastructure. It also aims to analyze various blockchain consensus mechanisms, with a particular focus on their energy consumption and feasibility for practical deployment. Based on these insights, the study seeks to design a blockchain-based system that emphasizes ease of integration, compatibility with existing legacy systems, and maintains acceptable levels of energy consumption. Finally, the study will simulate and compare the energy consumption of the proposed blockchain system against that of a traditional centralized model using CloudSim Plus.

This study presents a simulation-based feasibility analysis comparing classical and blockchain-based healthcare system architectures, using workload data from 2023. The focus is limited to execution-phase energy consumption, excluding idle power, cost, security, prototype development, or smart contract implementation. The aim is to inform decisions about blockchain adoption in healthcare environments.

2. Materials and Methods

2.1 Consensus mechanism

Consensus mechanisms underpin the integrity and reliability of blockchain systems by ensuring that all participating nodes agree on the data state. In healthcare systems, where operational sustainability and energy efficiency are crucial, the choice of consensus model must balance trust, fault tolerance, and computational cost. This section examines six widely studied consensus mechanisms—PoW, PoS, DPoS, BFT, PBFT, and PoA—through the lens of their energy demands and suitability for a national healthcare network. Proof of Work (PoW) is commonly used in permissionless blockchains such as Bitcoin. It relies on a competition-based process where nodes race to solve cryptographic puzzles by iterating nonces to produce a valid hash. This approach is secure but computationally expensive due to massive parallel processing and dynamically adjusted difficulty levels. As a result, PoW consumes extremely high amounts of electricity, making it one of the least energy-efficient consensus mechanisms. Its intensive energy profile renders it inappropriate for healthcare systems that require cost-effective and environmentally sustainable operation.

Proof of Stake (PoS) was introduced to mitigate PoW's energy consumption. Instead of brute-force computation, PoS selects validators based on their token holdings, significantly reducing computational demands. However, while PoS improves efficiency, it still requires continuous operation and selection algorithms that consume considerable energy, particularly in large networks. Additionally, PoS introduces centralization concerns by favoring wealth accumulation.

Delegated Proof of Stake (DPoS) optimizes PoS by reducing the validator pool to a fixed number of elected delegates. This substantially lowers energy use, as fewer nodes actively validate blocks. Yet, DPoS remains a permissioned system and still inherits centralization and transparency trade-offs. While energy consumption is reduced, governance integrity can be compromised.

Byzantine Fault Tolerance (BFT) and Practical Byzantine Fault Tolerance (PBFT) are consensus mechanisms designed to achieve agreement even in the presence of faulty or malicious nodes. BFT relies on a deterministic communication model that is energy-efficient in small, trusted networks but faces scalability issues due to quadratic communication overhead. PBFT enhances BFT by optimizing message exchanges and reducing latency, making it more practical for real-world, permissioned environments. However, both mechanisms require substantial communication between nodes, which can limit their performance in large-scale systems. They are most effective in controlled, small-to-medium networks where trust and participation are well-defined.

Proof of Authority (PoA) is recognized as the most energy-efficient consensus mechanism. Operating in a permissioned environment, PoA employs a small, pre-approved group of trusted validators. It eliminates probabilistic selection and computation-intensive competition entirely, which minimizes energy consumption. Though it reduces decentralization, this trade-off is aligned with healthcare governance structures, where institutional trust and regulatory oversight are already present.

As shown in Fig. 1, PoA demonstrates the lowest indicative energy consumption among all evaluated mechanisms. Its lightweight operational model and compatibility with regulated environments make it the optimal choice for implementing blockchain-based healthcare data sharing in Thailand.

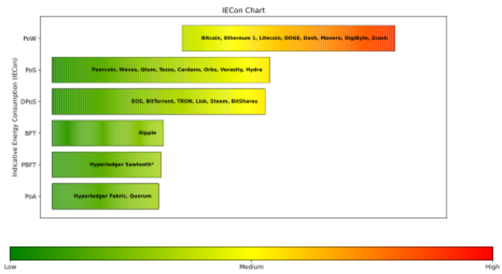

Fig. 1. Indicative energy consumption for each consensus.

2.2 CloudSim Plus

CloudSim Plus is an open-source simulation framework written in Java, specifically designed for modeling and evaluating cloud and edge computing environments. It employs an event-driven updating mechanism, allowing the simulation to advance only when specific events occur. This design enhances simulation efficiency and makes the framework particularly well-suited for precise energy consumption analysis in complex system architectures. The framework features a modular, object-oriented structure composed of essential components such as Datacenters, Hosts, Virtual Machines (VMs), Cloudlets, PowerModels, Utilization Models, and CloudletSchedulers. A Datacenter represents a logical collection of hosts; each Host simulates a physical server; and each VM executes computational tasks. Cloudlets represent individual workload units and are defined by their computational length in Million Instructions (MI). Hosts are characterized by their capacity in Million Instructions Per Second

(MIPS), enabling accurate modeling of execution times.

CloudSim Plus supports different types of CloudletSchedulers, which determine how VMs allocate CPU resources to cloudlets. Two primary scheduling policies are provided: SpaceShared, where each cloudlet has exclusive access to a processing element until completion, and TimeShared, where processing capacity is divided among cloudlets running concurrently. In this study, the TimeShared scheduler is used to simulate realistic multitasking and overlapping task execution as seen in healthcare systems.

To simulate energy consumption, PowerModels are assigned to hosts, estimating power usage based on real-time CPU activity. The power consumed at any given time is calculated using the equation:

$$Power = CPU\ Load \times Power\ Model, \tag{2.1}$$

where CPU Load is determined by the utilization level of the host, and Power Model reflects the corresponding energy profile derived from empirical data. Coupled with Utilization Models, which simulate workload variation over time, this configuration allows for dynamic and accurate energy profiling.

CloudSim Plus thus provides a flexible and reliable environment for simulating the performance and energy behavior of healthcare information systems across both traditional and blockchain-based infrastructures.

3. Proposed System Design

3.1 Thailand Healthcare architecture

Thailand's healthcare system is characterized not only by varying levels of IT readiness and institutional fragmentation but also by a well-established hierarchical data structure. In this model, hospitals and health units submit patient records to regional data centers, which then relay the data to national systems.

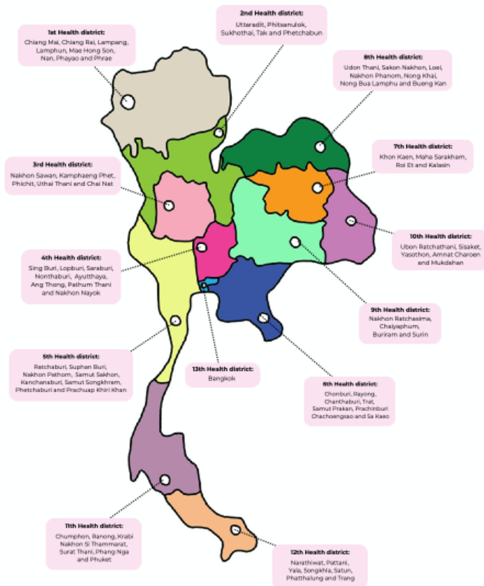

Fig. 2. Thai Healthcare District Structure.

This study designs its system to follow the existing structure, as shown in Fig. 2, for compatibility.

Moreover, the 13th health district—comprising only Bangkok—is treated uniquely due to its high population density and the advanced readiness of its medical centers. Bangkok is subdivided into seven health zones as shown in Fig. 3, which this study adopts as the basis for node assignment in the proposed blockchain system design.

3.2 Insurance organization

In Thailand, three main insurance organizations are included as trusted nodes in the blockchain system to enhance trans-

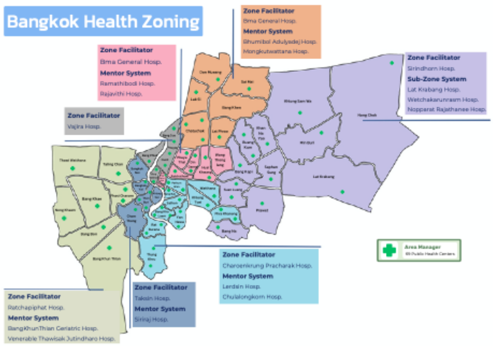

Fig. 3. Bangkok Health Zoning.

parency, traceability, and claims processing:

- National Health Security Office (NHSO): Manages the Universal Coverage Scheme (UCS).
- Office of Insurance Commission (OIC): Regulates and supervises private insurance providers.
- Social Security Office (SSO): Manages the Social Security Scheme (SSS) for formal-sector employees.

The blockchain system in this study was designed with a total of 23 nodes to reflect the organizational structure of Thailand's healthcare ecosystem. These include 12 regional health district nodes (excluding Health District 13, which overlaps with Bangkok), 7 Bangkok zoning nodes to account for the capital's population density and administrative segmentation, 3 insurance organization nodes representing SSO, OIC, and NHSO, and a master node operated by MOPH to oversee and coordinate national-level operations.

3.3 Energy factors

The energy consumption of blockchain-based healthcare applications is mainly influenced by two factors: the consensus mechanism and the number of nodes.

PoA was chosen as the most efficient consensus mechanism in a permissioned environment, relying on pre-selected validators and avoiding competitive selection, which reduces computational cost.

The number of nodes increases replication and inter-node communication overhead. Even with PoA, validators add to energy use; thus, maintaining a small, trusted validator pool improves efficiency.

3.4 Simulation testbed

3.4.1 General setting

The simulation aims to compare the energy consumption between a designed blockchain-based healthcare system and a classical centralized system using the CloudSim Plus framework. The blockchain system includes 23 nodes, consisting of 12 health district nodes excluding Bangkok, 7 Bangkok zoning nodes, 3 insurance organization nodes, and 1 master node representing MOPH. In contrast, the classical system comprises 13 nodes, each representing one of Thailand's health districts. Accordingly, the 13-node centralized model serves as the classical benchmark, while the 23-node distributed model serves as the blockchain benchmark. Both benchmarks were executed under identical CloudSim Plus configurations using workloads defined in Million Instructions (MI) and processed on hosts characterized by Million Instructions Per Second (MIPS) capacity.

The simulation was executed on a laptop equipped with an Apple M2 chip featuring an 8-core CPU running at 3.5 GHz. Each cloudlet in the simulation represents a healthcare transaction, and each datacenter corresponds to a blockchain or classical node. The simulation is configured such that each datacenter contains exactly one host, and each host is assigned one virtual machine (VM) to ensure consistent benchmarking across all nodes.

3.4.2 Host specification

Each simulated node in the study is modeled after the Dell PowerEdge XR11 server, equipped with an Intel Xeon Gold 6338N processor running at 2.20 GHz. This hardware configuration was selected due to its widespread use in enterprise and data center environments, along with its fully documented performance specifications from the Standard Performance Evaluation Corporation (SPEC). These specifications include 32 CPU cores / 64 threads, 128 GB of RAM, 10 Gbps bandwidth, and 480 GB of storage. Selecting this server model ensures that the simulation reflects real-world infrastructure with credible, standardized, and verifiable performance characteristics.

The host's processing capacity is estimated using the following formulas:

$$MIPS = \frac{2 \times FLOPS}{10^6}, \tag{3.1}$$

$$FLOPS = Core \times Clock \times IPC, \tag{3.2}$$

where IPC (Instructions Per Cycle) represents how many instructions a CPU can execute in a single clock cycle, and is derived by dividing the processor's Advanced Vector Extensions (AVX) width by the instruction width, based on the processor's vector architecture and instruction set configuration.

$$IPC = \frac{AVX\ width}{Instruction\ width}, \tag{3.3}$$

$$MIPS = \frac{2 \times Core \times Clock \times IPC}{10^6}, \tag{3.4}$$

Substituting the values:

$$MIPS = 2000 \times 32 \times 2.2 \times (\frac{512}{16})$$

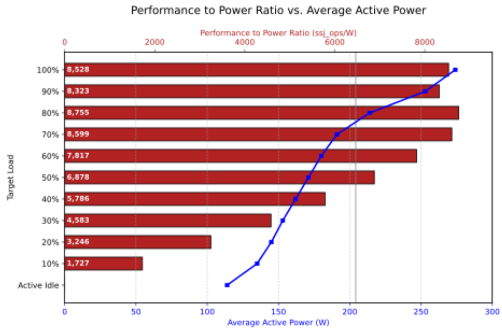

Fig. 4. Dell Inc. PowerEdge XR11 (Intel Xeon Gold 6338N, 2.20 GHz) power reference from The Standard Performance Evaluation Corporation (SPEC).

$$MIPS = 2,252,800$$

Each datacenter in the simulation is configured with one host and one virtual machine to ensure consistent benchmarking. Energy consumption is estimated using a PowerModel based on CPU utilization over time, with power characteristics derived from real hardware data (SPEC), as shown in Fig. 4 This setup enables realistic, workload-sensitive comparisons of energy efficiency between classical and blockchain-based architectures.

3.4.3 Workload configuration

The simulation workload is derived from 43 standard healthcare transaction files recorded monthly in 2023, reflecting actual operational data from Thailand's national healthcare system. The dataset includes approximately 75 million transactions per month, with an estimated monthly growth of 300,000 transactions. Given this enormous volume, the simulation is scaled down to represent 30 minutes of transaction activity from each month, preserving the statistical properties of real-world usage while ensuring computational feasibility.

Each transaction is modeled as a cloudlet, representing a discrete unit of computational work. Cloudlet arrivals follow a Poisson distribution, selected for its ability to capture the natural, independent, and irregular timing of transaction events. The distribution is governed by the probability mass function:

$$P(X = k) = \frac{e^{-\lambda} \cdot \lambda^k}{k!}, \qquad (3.5)$$

where λ is the average transaction arrival rate (in transactions per second), and k is the number of arrivals within a one-second interval. In this simulation, λ is calculated from the transactions-per-second (TPS) derived from the dataset.

To simulate Poisson-distributed transaction arrivals, this study uses the sample method from the Apache Commons Math library. This method automatically chooses between two techniques based on the average arrival rate. If the average is small—forty or less—it uses an inversion method that estimates how many events occur by multiplying random numbers until a certain limit is reached. For higher average rates, it switches to a rejection sampling method, which is more efficient under heavy loads. This approach generates sample values using a combination of normally distributed and exponentially distributed random numbers. By adapting the sampling technique to the expected arrival rate, the system can simulate healthcare transaction patterns efficiently and accurately.

The classical model is based on the actual transaction distribution, with each transaction executed in 0.5 seconds and an average CPU utilization of 70%.

The blockchain-based model uses a workload derived from the average transaction volume, evenly distributed across participating nodes. This model features a longer execution time of 2 seconds per transaction, which represents the additional processing required for transaction validation and encryption. It also operates with a higher CPU utilization of 90%, reflecting the computational overhead introduced by distributed consensus and inter-node communication.

4. Results and Discussion

4.1 43-file standard transactions

The 43-files standard is a national healthcare data schema mandated by Thailand's Ministry of Public Health. It aggregates standardized health transaction records from hospitals and public health units into regional data centers. Each of Thailand's 13 health districts submits monthly data to the Thai Health Data Center (HDC), forming the backbone of national healthcare analytics, monitoring, and reporting.

To estimate future workload and support realistic transaction simulation, historical data spanning 2018 to 2023 was analyzed. A linear regression was applied to the total volume of 43-file transactions aggregated across all 13 health districts. The resulting regression model, as shown in Fig. 5, is:

$$y = 304361.58x + 74216552.11 \qquad (4.1)$$

where y represents the estimated number of transactions per month, and x denotes the number of months since the starting point of the dataset.

This model enables the projection of future transaction volumes, supporting workload scaling and forecasting in healthcare information systems.

4.2 Classical simulation result

In Fig. 6, the x-axis represents the months of the year 2023, while the y-axis shows energy consumption in kilowatt-

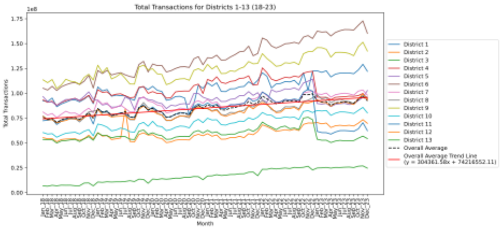

Fig. 5. Transaction growth of 43-file standard data across Thailand's 13 health districts from 2018 to 2023, with linear regression trendline.

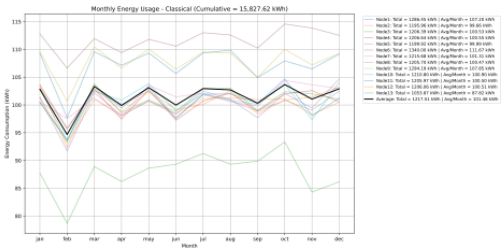

Fig. 6. Energy consumption in the classical system.

hours (kWh); each line represents one of the 13 health district nodes. Energy consumption remained generally stable throughout the year, with fluctuations closely aligned with the number of days in each month. Higher consumption was observed in longer months (e.g., January, March, May), while shorter months like February showed lower usage. These variations reflect the direct impact of calendar length on total workload. The simulation recorded a total annual energy consumption of 15,827.62 kWh, with an average monthly per-node consumption of 101.46 kWh and a total per-node annual average of 1,217.51 kWh, demonstrating the efficiency and predictability of the centralized model.

4.3 Blockchain simulation result

Fig. 7, similar to Fig. 6, illustrates energy consumption within the blockchain system. It features 23 distinct colored lines representing the 23 nodes in the blockchain

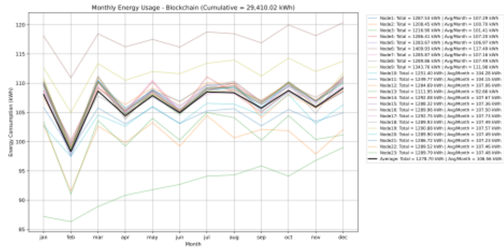

Fig. 7. Energy consumption in the blockchain system.

network, all following a trend comparable to the classical system simulation. The simulation reported a total annual energy consumption of 29,410.02 kWh, with an average per-node annual consumption of 1,278.70 kWh. The average monthly energy consumption per node was 106.56 kWh, which is slightly higher than that observed in the classical system, reflecting the additional overhead introduced by the decentralized architecture.

4.4 Comparison result

The energy comparison between the classical and blockchain systems reveals a significant difference of 13,582.4 kWh annually, representing a 1.86-fold increase in consumption for the blockchain model, as illustrated in Fig. 8. However, this additional energy cost brings notable functional benefits. The blockchain system helps distribute workload more evenly, reducing pressure on central components and mitigating the risk of single points of failure. Its inherent immutability also enhances data integrity. Together, these features contribute to a more robust and adaptable infrastructure, better suited to support the future needs of national healthcare systems.

5. Conclusion

This study proposes a blockchain-based architecture tailored for practical de-

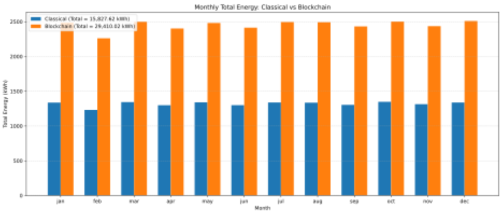

Fig. 8. Energy consumption comparison between classical and blockchain systems.

ployment in Thailand's healthcare network, focusing on secure, resilient, and continuous data exchange. A major concern in blockchain adoption is its high energy consumption. This research identifies two main contributors to this issue: the consensus mechanism and the number of participating nodes. Traditional consensus methods like PoW are computationally intensive and unsuitable for energy-conscious public health applications. To address this, the study adopts PoA—a consensus mechanism optimized for private, permissioned blockchains. PoA enhances energy efficiency by eliminating competitive validator selection, instead using pre-identified, trusted validators. This approach significantly reduces computational overhead while maintaining security and performance. The system design includes 23 authorized nodes, reflecting an area-based management strategy: 12 of the 13 health districts (excluding District 13, Bangkok), replaced by 7 Bangkok health zones due to urban density, 3 insurance-related nodes (SSS, OIC, NHSO), and a master node under the Ministry of Public Health (MOPH).

Energy consumption was evaluated using the CloudSim Plus simulator, comparing this blockchain system to a traditional centralized model with 13 health district data centers. The simulation results indicate an annual execution energy consumption of 13,582.4 kWh for the blockchain-based model, representing a 1.86-fold increase compared to the classical system. The classical system used 15,827.62 kWh/year, while the blockchain system consumed 29,410.02 kWh/year. Despite higher consumption, the blockchain model offers superior data integrity, availability, and system resilience, making it a practical and strategic solution for national healthcare modernization. This finding is expected to support decision-making regarding blockchain adoption within the Thai healthcare network and guide future studies on its feasibility and sustainability.

To enable more precise simulation and evaluation, future research could leverage higher-specification computational resources to model dynamic, real-world workloads. Instead of the current static, scaled-down approach (30 minutes per month), simulations should reflect continuous operations, including peak hours (e.g., around 10:00 AM), off-peak periods at night, and differing activity levels between weekdays and weekends. Incorporating idle-state energy consumption would provide a fuller picture of total system demands. Furthermore, integrating key management and data encryption flows is essential for evaluating security, privacy, and cryptographic overhead in practical blockchain-based healthcare systems. In addition, while this study relies on the 43-file standard transactions healthcare dataset as the independent variable for workload simulation, future research should validate the findings with additional datasets such as hospital admission records, claims data, or electronic medical records to ensure broader applicability.

Acknowledgements

The author would like to express sincere gratitude to the Department of Electrical and Computer Engineering, Thammasat School of Engineering, Thammasat University, for the valuable guidance, constructive feedback, and continuous support provided throughout this research. Appreciation is also extended to colleagues and the department for offering technical resources and fostering an encouraging research environment.

References

[1] NHSO Regional Office [Internet]. Nhso.go.th. 2024 [cited 2025 Apr 21]. Available from: https://eng.nhso.go.th/view/1/NHSO_Regional_Office/EN-US

[2] Lertpitakpong C. Strengthening health information system on hypertension: 43 folders. PH.PL [Internet]. 2021 Nov 19 [cited 2025 Apr 21];7(3):573–82. Available from: https://so05.tci-thaijo.org/index.php/journal_law/article/-view/255699

[3] Quattrocchi G, Scaramuzza F, Tamburri D. The Blockchain Trilemma: An Evaluation Framework. IEEE Software. 2024;PP:1–8.

[4] Bashir I. Mastering blockchain: distributed ledgers, decentralization and smart contracts explained. Birmingham (UK); Mumbai (India): Packt; 2017.

[5] Solat S, Calvez P, Nait-Abdesselam F. Permissioned vs. permissionless blockchain: how and why there is only one right choice. J Softw. 2020;16:95–106.

[6] Cole R, Cheng L. Modeling the energy consumption of blockchain consensus algorithms. In: Proceedings of the 2018 IEEE International Conference on Internet of Things (iThings), IEEE GreenCom, IEEE CPSCom, and IEEE SmartData; 2018 Jul 30–Aug 3; Halifax, Canada. Piscataway (NJ): IEEE; 2018. p. 1691–6.

[7] Sedlmeir J, Buhl HU, Fridgen G, Keller R. The energy consumption of blockchain technology: beyond myth. Bus Inf Syst Eng. 2020;62(6):599–608.

[8] Kohli V, Chakravarty S, Chamola V, Sangwan KS, Zeadally S. An analysis of energy consumption and carbon footprints of cryptocurrencies and possible solutions. Digit Commun Netw. 2022;9(1):79–89.

[9] Bada AO, Damianou A, Angelopoulos CM, Katos V. Towards a green blockchain: a review of consensus mechanisms and their energy consumption. In: Proceedings of the 2021 17th International Conference on Distributed Computing in Sensor Systems (DCOSS); 2021 Jun 7–9; Pafos, Cyprus. Piscataway (NJ): IEEE; 2021. p. 503–11.

[10] Silva Filho MC, Oliveira RL, Monteiro CC, Inacio PRM, Freire MM. CloudSim Plus: a cloud computing simulation framework pursuing software engineering principles for improved modularity, extensibility and correctness. In: Proceedings of the 2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM); 2017 May 8–12; Lisbon, Portugal. Piscataway (NJ): IEEE; 2017. p. 400–6.